Federated Learning for E-Commerce Recommendation Systems: A Survey on Privacy-Preserving Personalization

Nier Mou

Department of Electrical and Electronic Engineering Department, Hong Kong Polytechnic University, Hong Kong, China

24110607d@connect.polyu.hk

Abstract. In the information age, while e-commerce recommendation systems enhance user experience, they also face a significant risk of user privacy leakage. Traditional privacy protection technologies (such as K-anonymity and differential privacy) often come at the expense of data utility and recommendation accuracy. Federated learning, as an emerging distributed machine learning paradigm, offers a promising solution to the tension between privacy and personalization through its core principle of 'data remains local, models move". This paper reviews the latest progress in federated learning for e-commerce recommendation, focusing on analyzing federated recommendation architectures based on deep neural networks (such as FedPrivRec) and natural language processing, as well as their applications in cross-merchant collaboration, privacy protection, and other areas. Meanwhile, this paper examines the primary challenges currently faced by federated recommendation systems, including privacy attacks such as member inference, data heterogeneity, as well as communication and computational efficiency bottlenecks and cost pressures. Finally, future research directions, such as the optimization of communication and hardware computing power, were prospected, aiming to serve as a reference for building a privacy-secure, efficient, and sustainable e-commerce recommendation system service.

Keywords: Federated learning; data privacy; recommendation system; E-Commerce.

1. Introduction

In the information age, the rapid increase of online activities, such as shopping, socializing, and entertainment, has generated a vast amount of behavioral data. Theoretically, this data deluge provides an ideal foundation for artificial intelligence systems to understand and predict human preferences with ever-increasing accuracy. However, rising risks of personal data leakage and growing awareness of privacy risks have jointly led to a decline in users' trust in the platform providing network services. Most users were increasingly unwilling to share their data, so platforms started to adopt various privacy-preserving techniques to protect user privacy, such as k-anonymity, data masking, differential privacy, cryptographic techniques, and data minimization. Without exception, all these mature privacy-enhancing technologies share a similar problem that perturbation to the data will irreversibly reduce the granularity and dimension of the data. Grigorios Loukides and Jianhua Shao demonstrate that k-anonymity can cause loss of significant information from the dataset, because data generalization reduces the precision and level of detail in the data [1]. A similar drawback is observed in data masking techniques [2]. Differential privacy achieves privacy protection by adding noise to data. The amount of noise is directly proportional to security intensity but inversely proportional to the accuracy of the data [3]. These general techniques are commonly valuable for reducing the risks of identifying individuals, frequently undermine the value and utility of user data.

The tension between learning customers' preferences and ensuring customers' privacy puts the platforms that need to apply a recommender system in a dilemma. Commercially, they must enhance the accuracy of recommended content to improve user engagement. Meanwhile, platforms must also comply with regulatory requirements under frameworks such as the GDPR and CCPA, which enforce principles like explicit user consent and data localization. The e-commerce industry is one of the fields that was significantly influenced by this conflict. In recent years, the application of federated learning (FL) is regarded as a promising technology in the e-commerce domain. FL

does not collect raw data to a central server but operates by distributing the global model to local nodes on the device for training. It directly solves users' worries about privacy data leakage, because data will not even leave their electronic devices. Also, benefiting from this technology, e-commerce platforms can not only obtain effective user information to the greatest extent but also are able to achieve the exchange and sharing of learning outcomes with other platforms for collaboration. Achieving the sharing and circulation of data between different data sources breaks down data islands and promotes sustainable development in e-commerce [4]. This review will demonstrate the latest evolution of federated learning architectures designed for recommendation systems, evaluate their efficacy in overcoming the privacy-personalization paradox, and discuss the persistent challenges and future directions shaping this critical field of research.

2. The Application of FL in E-commerce

2.1. Introduction of FL

Federated Learning (FL) is a decentralized machine learning approach that brings learning to the edge or directly on the device. As shown in [5], FL avoids collecting raw data from distributed sources by sending an initial global model to each participating client, such as smartphones, IOT devices, and local servers. Each Client trains the model using local data and updates the model locally, including gradients and weights. Only the updated model will be transmitted back to the central server. These updates are securely aggregated to promote a global model, which will be redistributed. During the whole process, clients do not have data leakage risks. FL fundamentally enhances privacy and security through architecture designs, observing strict regulations like GDPR and CCPA. Compared with traditional Machine Learning (ML) techniques, FL is especially valuable in scenarios involving sensitive and decentralized data, such as personalized recommendation systems, healthcare analytics, and industrial IOT. FL can generally be classified as Horizontal Federated Learning (HFL), Vertical Federated Learning (VFL), and Federated Transfer Leaning (FTL), and their application scenarios are different [6]. HFL promotes similar organizations to commercially collaborate while protecting respective commercial data security. VFL is more suitable for cross-industry cooperation. For example, an e-commerce platform can cooperate with an e-payment platform to jointly train a malicious order rigging identification model, with the former providing behavioral characteristics and the latter providing financial characteristics. Using FTL can quickly launch blank products or market projects.

2.2. FedPrivRec and PrivRec Based on DNN

Federated Privacy-Recommendation model (FedPrivRec) shown in Fig. 1 is a novel federated learning framework that is specifically designed for real-time recommendation while protecting privacy in e-commerce contexts. It adopts a hierarchical federated aggregation strategy, including 4 main layers in this framework: client devices, edge aggregators, central coordinator, and model repository [7]. The following flowchart demonstrates each layer of FedPrivRec and data flow. The red arrow indicates the upward update of the encryption model parameters, and the blue arrow indicates the downward distribution of the encryption model.

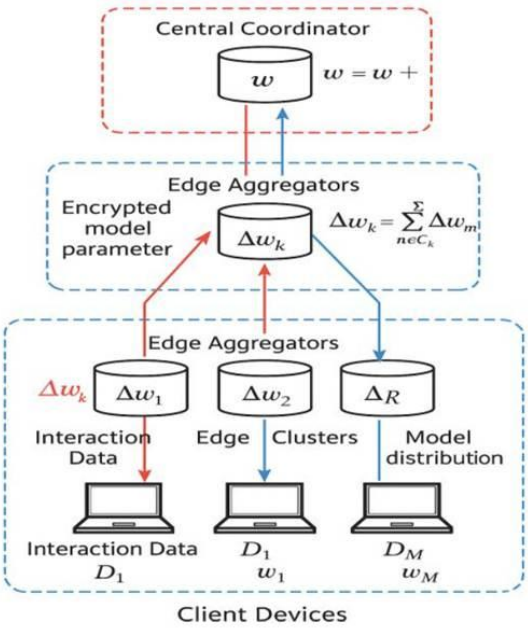


Fig. 1 A workflow of the common FedPrivRec model [7]

Based on the common FedPrivRec structure keeping sensitive user data localized, Wang et.al. [8], incorporates a deep neural network (DNN) with basic FedPrivRec. The core capability of DNN is to map high-dimensional, sparse raw features (some features used in the e-commerce scene, such as the user's age, occupation, item category, and text description) to a low-dimensional, dense vector space (i.e., embedding vectors), which avoids reliance on unique user identifiers and further enhances the safety level of local data. In the framework of combining deep neural networks (DNN) with federated learning (FL), the model structure consists of a central server and multiple distributed clients. Each client uses local data to train the same DNN model. The workflow begins with the server initializing the global DNN model and distributing it to each client. Subsequently, each client performs multiple rounds of local training, calculating model gradients or weight updates. After completion, the clients only upload the model updates rather than the original data encrypted to the server. Finally, the server integrates all the updates through a secure aggregation algorithm (such as FedAvg) to generate a new generation of the global model and re-distribute it. This process repeats until the model converges. This model can achieve collaborative training of high-performance DNN models while fully protecting the original data of users.

2.3. NLP-based Recommendation System Employing Transformer-Based Architecture

Another research also makes a creative technical combination based on FL constructure shown in Fig. 2, focusing on enhancing the flexibility of recommendation and application of multi-merchant e-commerce scenarios [9]. This tokenization strategy, combined with federated learning, allows for semantic understanding without exposing raw text data. To overcome the reliance of traditional recommendation systems on a single product ID, researchers have adopted a tokenization strategy based on spatial word segmentation. The names of goods and merchants are decomposed into discrete semantic units (tokens). For example, "Seafood Pizza" from "Pizza Hut" is transformed into ['Seafood', 'Pizza', 'Pizza Hut']. Tokenization enables the model to understand the semantic association between "Seafood Pizza" and "Seafood Pasta", even if they come from different merchants, improving the flexibility of recommendation. Based on the semantic parsing of product names and merchant names through the tokenization strategy, they further designed a hierarchical parameter sharing mechanism to address the current situation where the data volume and vocabulary distribution among different merchants are extremely unbalanced.

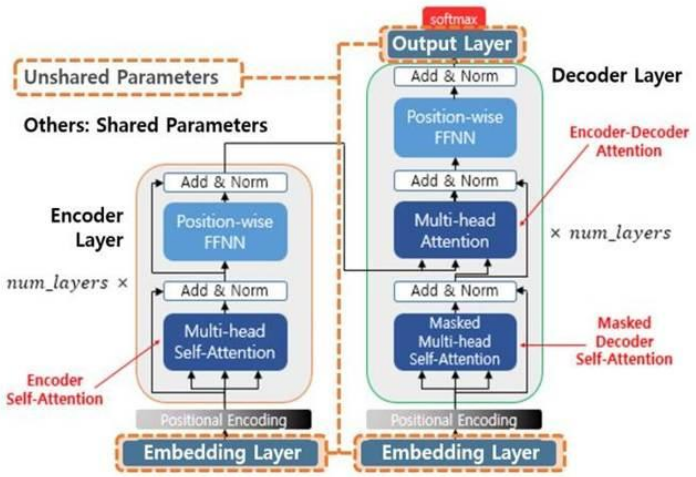


Fig. 2 FL of Transformer with Glocalization [9]

This mechanism sets that only the intermediate layers of the Transformer encoder and decoder participate in the federated averaging (FedAvg) aggregation, while the Embedding Layer and Output Layer are completely retained locally and do not participate in global sharing. This design is mainly based on two core considerations:

First, the independence of the vocabulary. The localized embedding layer enables each client to flexibly adapt to its own product directory and feature space.

Second, the balance between knowledge sharing and privacy protection. The intermediate layers learn abstract semantic features (such as "food type"), and their federated aggregation can effectively achieve knowledge transfer and model enhancement. The embedding layer contains sensitive mapping information of the original vocabulary. Strictly retaining the embedding layer locally not only ensures the autonomy of merchant data but also eliminates the possibility of privacy vocabulary information leakage at the source.

3. Discussion

3.1. Challenges

Various technologies related to federated learning can help e-commerce enterprises and platforms reduce risks of data leakage from technical structure improvement, while flexibly adapting to different recommendation scenarios. However, after embedding FL, platforms and enterprises still need to cope with possible malicious attacks. Several common attack sources include malicious platforms, malicious tenants, and malicious users. The final trained model usually incurs three types of privacy attacks: parameter inference, model inversion, and membership inference. The basic design of FL is that raw data never leaves local devices, but parameter information is deliberately shared, and model inversion requires auxiliary knowledge about specific user data, which is inaccessible to attackers in the FL environment. Thus, membership inference becomes the main attack approach. Melis et.al found that malicious actors (for example, curious servers or clients) can leverage exposed model gradients not only to perform member inference but also to infer sensitive global attributes of local datasets, and even uniquely re-identify individual clients [10]. To defend against this attack, DF was proposed to be applied in the FL structure [8]. DF adds some additional noises in gradients to make it harder for attackers to reliably infer the information of any single training sample from the released model or update. However, DP also introduces a trade-off: the addition of noise inevitably degrades recommendation performance. Focusing on this problem, Melis et.al also propose a scheme to reduce the impact of noise through two-stage training: Firstly, high-quality item representations are trained through self-supervised learning without noise intervention; Subsequently, these strong representations are taken as

initialization parameters to conduct user-item interaction modeling under differential privacy protection, thereby reducing the negative impact of noise on model performance.

The complexity of user types, product types, cross-temporal and spatial features determines that e-commence field have to face the problem of data heterogeneity. Data heterogeneity reduces the possibility that all clients simultaneously accept a common collaboration model and damages the training efficiency and accuracy of the federated learning model serving recommendation systems [11]. Liu et.al present that FPFC [12], a new clustering method based on nonconvex pair fusion penalty, can solve the problem of global model divergence caused by data heterogeneity by punishing the model according to the difference between models to promote the model differences to 0 to cluster data. Another more flexible approach is proposed by Panchal et.al [13]. They found that in some clients, using a global model that is trained by data from clients to establish a personalized model for recommendation is possible to achieve lower accuracy than the model without personalization. Since the global model often fails to achieve the expected accuracy and the technology for solving data heterogeneity is complex, they proposed a per-instance personalization FL algorithm, Flow. Flow can establish dynamic personalized models that could determine more effective and efficient decisions by evaluating the preference between local parameters and global parameters. This solution ingeniously avoids the high time and computing power consumption faced by large-scale data heterogeneity.

3.2. Prospects

FL recommendation system is a promising and burgeoning technique in the e-commerce domain, but some limitations still lack technical approaches. Although its outstanding efficacy in privacy security has been confirmed in academia, the real-world deployment of FL-based recommendation systems is still in its early stages. Except for some problems that have had some effective solutions involved in previous text, such as attack risks and data heterogeneity, there are still some problems that have technical or hardware bottlenecks.

Communication cost and efficiency are one of the biggest bottlenecks for FL. Due to privacy worries for the transmission of raw data, each device will definitely save the generated data locally, so federated networks may have plenty of devices, which also results in limitations on the available bandwidth, energy, and power in the federated networks [14]. Therefore, methods with high communication efficiency are required. These methods are better to have as few rounds of communication as possible and as small a message size as possible so that communication in such a setup can be further reduced. In the future, one of the feasible approaches is to integrate 5G or 6G with FL, capitalizing on their benefits, such as high speed, low latency, large connection volume, and flexible openness to solve this problem [11].

In addition to the network’s aspect, computing hardware is another critical factor that needs future development. Current popular e-commerce platforms such as Amazon and Taobao have billions of users. Meanwhile, FL involves complex aggregation algorithms and requires frequent model updates. Such a huge number of users will cause tremendous pressure on computing and storage. If the adoption of FL recommendation systems becomes a trend in the e-commerce industry, addressing their high threshold for server computing power should be the first step.

4. Conclusion

Federated learning, through its distributed architecture design, offers a revolutionary technical path for the e-commerce industry to achieve efficient personalized recommendations and break down data islands under strict data privacy regulations. This article first reviews the basic principles of federated learning and its diverse applications in the e-commerce field. Subsequently, it demonstrated its integration with technologies such as deep learning and natural language processing. After the integration of technologies, precise product recommendations can still be achieved while ensuring that users' original data remains local.

However, federated learning recommendation system for large-scale, mature commercial use still faces many challenges. Firstly, privacy and security threats, especially member inference attacks, require the system to integrate enhanced technologies such as differential privacy, but this will introduce noise and lead to a decline in recommendation performance. Secondly, the heterogeneity of user and product data in the e-commerce scenario poses a severe test to the generalization ability of the global model. Future research work should develop more advanced personalized federated learning algorithms. Finally, the huge communication overhead and the heavy computing and coordination pressure on the central server constitute a major system bottleneck.

Future research will need interdisciplinary efforts to focus on: developing aggregation algorithms with higher communication efficiency and stronger privacy protection; exploring the integration with advanced network technologies such as 5G/6G to reduce latency; designing personalized solutions that can effectively deal with data heterogeneity; and building a powerful computing infrastructure capable of supporting a user scale of hundreds of millions. Only by overcoming these obstacles can federated learning truly unleash its potential and become a key driving force for the sustainable development of the e-commerce industry.

References

[1] Loukides G, Shao J. Capturing data usefulness and privacy protection in K-anonymisation. In: Proceedings of the 2007 ACM Symposium on Applied Computing; 2007 Mar 11–14; Seoul, Korea. New York: ACM; 2007. p. 370–4.

[2] Duncan G, Stokes L. Data masking for disclosure limitation. Wiley Interdisciplinary Reviews: Computational Statistics, 2009, 1(1): 83-92.

[3] Zhao J, Chen Y, Zhang W. Differential privacy preservation in deep learning: Challenges, opportunities and solutions. IEEE Access, 2019, 7: 48901-48911.

[4] Li J T, Cui T X, Yang K W, et al. Demand forecasting of E-commerce enterprises based on horizontal federated learning from the perspective of sustainable development. Sustainability, 2021, 13(23): 13050.

[5] Mothukuri V, Parizi R M, Pouriyeh S, et al. A survey on security and privacy of federated learning. Future Generation Computer Systems, 2021, 115: 619-640.

[6] Yang Q, Liu Y, Chen T, et al. Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems and Technology, 2019, 10(2): 1-19.

[7] Ji S, Liang Y, Xiao X, et al. An attitude-adaptation negotiation strategy in electronic market environments. In Proceedings of the Eighth ACIS International Conference on Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing, Qingdao, China, 2007, 3: 125-130.

[8] Wang Q Y, Yin H Z, Chen T, et al. Fast-adapting and privacy-preserving federated recommender system. The VLDB Journal, 2022, 31(5): 877-896.

[9] Lee K J, Jeong B, Kim S, et al. General commerce intelligence: Glocally federated NLP-based engine for privacy-preserving and sustainable personalized services of multi-merchants. In Proceedings of the AAAI Conference on Artificial Intelligence, Vancouver, Canada, 2024, 38(21): 22752-22760.

[10] Melis L, Song C, De Cristofaro E, et al. Exploiting unintended feature leakage in collaborative learning. In Proceedings of the IEEE Symposium on Security and Privacy, San Francisco, CA, USA, 2019: 691-706.

[11] Pei J, Liu W, Li J, et al. A review of federated learning methods in heterogeneous scenarios. IEEE Transactions on Consumer Electronics, 2024, 70(3): 5983-5999.

[12] Yu X, Liu Z, Sun Y, et al. Clustered federated learning for heterogeneous data (student abstract). In Proceedings of the AAAI Conference on Artificial Intelligence, Washington, DC, USA, 2023, 37: 16378-16379.

[13] Panchal K, Choudhary S, Parikh N, et al. Flow: per-instance personalized federated learning. Advances in Neural Information Processing Systems, 2023, 36: 18712-18755.

[14] Bharati S, Mondal M R H, Podder P, et al. Federated learning: Applications, challenges and future directions. International Journal of Hybrid Intelligent Systems, 2022, 18(1-2): 19-35.